

EtherSAM: THE NEW STANDARD IN ETHERNET SERVICE TESTING

By Thierno Diallo and Marquis Dorais, Product Specialists, Transport and Datacom Business Unit

INTRODUCTION

With Ethernet continuing to evolve as the transport technology of choice, networks have shifted their focus from purely moving data to providing entertainment and new applications in the interconnected world. Ethernet-based services such as mobile backhaul, business and wholesale services need to carry a variety of applications, namely voice, video, e-mail, online trading and others. These latest applications impose new requirements on network performance, and on the methodologies used to validate the performance of these Ethernet services.

This application note examines EtherSAM or ITU-T Y.1564, the new ITU-T standard for turning up, installing and troubleshooting Ethernet-based services. EtherSAM is the only standard test methodology that allows for complete validation of Ethernet service-level agreements (SLAs) in a single, significantly faster test, and with the highest level of accuracy.

1. THE REALITY OF TODAY'S NETWORKS

Ethernet networks are now servicing real-time and sensitive services. By service, we are referring to the various types of traffic that the network can carry. Generally, all network traffic can be classified under three traffic types: best effort, real-time and high-priority. Each traffic type is affected differently by the network characteristics and must be groomed and shaped to meet their minimum performance objectives.

Traffic Type	Main Applications	Examples of Services
Best-Effort Data	Non real-time or data transport	<ul style="list-style-type: none"> - Data - Internet access - FTP download /upload - Server, storage applications
Real-Time Data	Real-time broadcast that cannot be recreated once lost	<ul style="list-style-type: none"> - VoIP - IPTV, video on demand - Internet radio, TV - Internet gaming - Videoconference
High-Priority Data	Mandatory traffic used to maintain stability in the network	<ul style="list-style-type: none"> - OAM frames - Switching/routing control frames - Network synchronization such as SyncE, 1588v2

Table 1. Network traffic types

To assure quality of service (QoS), providers need to properly configure their networks to define how the traffic inside will be prioritized. This is accomplished by assigning different levels of priority to each type for service and accurately configuring network prioritization algorithms. QoS enforcement refers to the method used to differentiate the traffic of various services via specific fields in the frames, thus prioritizing frames for certain services over other frames. These fields make it possible for a network element to discriminate between and service high- and low-priority traffic.

2. IMPORTANCE OF SLA

A service-level agreement (SLA) is a binding contract between a service provider and a customer, which guarantees the minimum performance that will be assured for the services provided. These SLAs specify the key forwarding characteristics and the minimum performance guaranteed for each characteristic.

Key Performance Indicators	Best-Effort Data (Internet Access)	Real-Time Data	High-Priority Data
CIR (Mbit/s) (green traffic)	2.5	5	10
EIR (Mbit/s) (yellow traffic)	5	0	5
Frame delay (ms)	<30	<5	5-15
Frame delay variaton (ms)	n/a	<1	n/a
Frame loss (%)	<0.05	<0.001	<0.05
VLAN	300	100	200

Table 2. Key performance indicators (KPIs) for various traffic types

Customer traffic is classified into three traffic classes, and each is assigned a specific color: green for committed traffic, yellow for excess traffic and red for discarded traffic.

- > Committed information rate (CIR), or green traffic: refers to bandwidth that is guaranteed available at all times for a specific service; for green traffic, minimum performance objectives (i.e., key performance indicators or KPIs) are guaranteed to be met.
- > Excess information rate (EIR), or yellow traffic: refers to excess bandwidth above CIR that may be available depending on network loading and usage; for yellow traffic, minimum performance objectives are not guaranteed to be met.
- > Discarded, or red traffic: refers to traffic that is above the CIR or the CIR/EIR rate, and that cannot be forwarded without disrupting other services; red traffic is therefore discarded.

Traffic Class	Bandwidth	Performance Objective	KPI Level
Green traffic	0 to CIR	Guaranteed forwarding	KPIs are guaranteed
Yellow traffic	CIR to EIR	Best effort	KPIs are not guaranteed
Red traffic	> EIR or CIR	Discarded traffic	Not applicable

Table 3. Traffic Classes

2.1 KEY PERFORMANCE INDICATORS

KPIs are specific traffic characteristics that indicate the minimum performance of a particular traffic profile. Under green traffic condition, the network must guarantee that these minimum performance requirements are met for all forwarded traffic. Typical KPIs include:

2.1.1 Bandwidth

Bandwidth refers to the maximum amount of data that can be forwarded. This measurement is a ratio of the total amount of traffic forwarded during a measurement window of one second. Bandwidth can either be “committed” or “excess”, with different performance guarantees.

Bandwidth must be controlled, because multiple services typically share a link. Therefore, each service must be limited to avoid affecting another service. Generating traffic over the bandwidth limit usually leads to frame buffering, congestion and frame loss or service outages.

2.1.2 Frame Delay (Latency)

Frame delay or latency, is a measurement of the time delay between a packet's transmission and its reception. Typically, this is a round-trip measurement, meaning that it simultaneously calculates both the near-end to far-end, and the far-end to near-end directions. This measurement is critical for voice applications, as too much latency can affect call quality, leading to the perception of echoes, incoherent conversations or even dropped calls.

2.1.3 Frame Loss

Frame loss can occur for numerous reasons, such as transmission errors or network congestion. Errors due to a physical phenomenon can occur during frame transmission, resulting in frames being discarded by networking devices such as switches and routers based on the frame check sequence field comparison. Network congestion also causes frames to be discarded, because networking devices must drop frames to avoid saturating a link in congestion conditions.

2.1.4 Frame Delay Variation (Packet Jitter)

Frame delay variation, or packet jitter, refers to the variability in arrival time between packet deliveries. As packets travel through a network, they are often queued and sent in bursts to the next hop. Random prioritization may occur, resulting in packet transmission at random rates. Packets are therefore received at irregular intervals. This jitter translates into stress on the receiving buffers of the end nodes, which may result in buffers becoming overused or underused in the event of large swings of jitter.

Real-time applications such as voice and video are especially sensitive to packet jitter. Buffers are designed to store a certain quantity of video or voice packets, which are then processed at regular intervals to provide a smooth and error-free transmission to the end user. Too much jitter will affect the quality of experience (QoE), because packets arriving at a fast rate will cause the buffers to overfill, resulting in packet loss. Packets arriving at a slow rate will cause buffers to empty, leading to still images or sound.

3 CURRENT TESTING METHODOLOGY: RFC 2544

RFC 2544 has been the most widely used Ethernet service testing methodology. This series of subtests provides a methodology to measure throughput, round-trip latency, burst and frame loss.

It was initially introduced as a benchmarking methodology for network interconnect devices in the lab. However, since RFC 2544 was able to measure throughput, burstability, frame loss and latency, and because it was the only existing standardized testing methodology, it was also used for Ethernet service testing in the field.

While this testing methodology provides key parameters to qualify the network, it is no longer sufficient in terms of fully validating today's Ethernet services. More specifically, RFC 2544 does not include all required measurements, such as packet jitter, QoS measurement and multiple concurrent service levels. Additionally, since RFC 2544 requires the performance of multiple, sequential tests to validate complete SLAs, this test method takes several hours, proving to be both time consuming and costly for operators. It is now a requirement to simulate all types of services that will run on the network and simultaneously qualify all key SLA parameters for each of these services.

4 REVOLUTIONARY TESTING METHODOLOGY: ETHERSAM (ITU-T Y.1564)

To resolve issues with existing methodologies, the ITU-T has introduced a new test standard: the ITU-T Y.1564, which is aligned with the requirements of today's Ethernet services. EXFO was the first to implement EtherSAM—the Ethernet service testing methodology based on this new standard—into its Ethernet testing products.

EtherSAM enables complete validation of all SLA parameters within a single test to ensure optimized QoS. Contrary to other methodologies, it supports new multiservice offerings. In fact, EtherSAM can simulate all types of services that will run on the network, and simultaneously qualify all key SLA parameters for each of these services. It also validates the QoS mechanisms provisioned in the network to prioritize the different service types, resulting in more accurate validation, and much faster deployment and troubleshooting. Moreover, EtherSAM offers additional capabilities, such as bidirectional measurements.

EtherSAM (ITU-T Y.1564) is based on the principle that the majority of service issues are found in two distinct categories: a) in the configuration of the network elements that carry the service, or b) in the performance of the network during high load conditions when multiple services cause congestion.

4.1 SERVICE CONFIGURATION

Forwarding devices such as switches, routers, bridges and network interface units are the foundation of any network because they interconnect segments. These forwarding devices must be properly configured to ensure that traffic is adequately groomed and forwarded in accordance with its service level agreement.

If a service is not correctly configured on a single device within the end-to-end path, network performance can be greatly affected. This may lead to service outage and network-wide issues, such as congestion and link failures. Therefore, a very important part of the testing effort is to ensure that devices are properly configured and able to handle the network traffic as intended.

4.2 SERVICE PERFORMANCE

Service performance refers to the ability of the network to carry multiple services at their maximum guaranteed rate without any degradation in performance; i.e., KPIs must remain within an acceptable range.

As network devices come under load, they must make quality decisions. They must prioritize one traffic flow over another in order to meet the KPIs of each traffic class. This is necessary because as the amount of traffic flow increases, so does the likelihood of performance failures.

Service performance assessment must be conducted over a medium- to long-term period, because problems typically occur in the long term and will probably not be seen with short-term testing.

The focus of EtherSAM (ITU-T Y.1564) is therefore threefold:

- First, the methodology serves as a validation tool. It ensures that the network complies with the SLA by ensuring that a service meets its KPI performance objectives at different rates, and within the committed range.
- Second, the methodology ensures that all services carried by the network meet their KPI performance objectives at their maximum committed rate. This proves that under maximum load, the network devices and paths are able to service all the traffic as designed.
- Third, service performance testing can be performed over medium- to long-term test periods to confirm that network elements are able to properly carry all services while under stress during a soaking period.

5 ETHERSAM: TESTS AND SUBTESTS

EtherSAM is comprised of two tests: the service configuration test and the service performance test. The frame size used for the service configuration test and the service performance test can be constant, or a distribution of multiple frame sizes. The ITU-T Y.1564 has defined a variable frame-size sequence format named EMIX, or Ethernet Mix. The EMIX frame sequence format can be configured from two to eight frames, with configurable frame sizes ranging from 64 to 16000 bytes. EMIX's main purpose is to emulate real-life network traffic and uncover potential issues that may not arise when testing with a constant frame size.

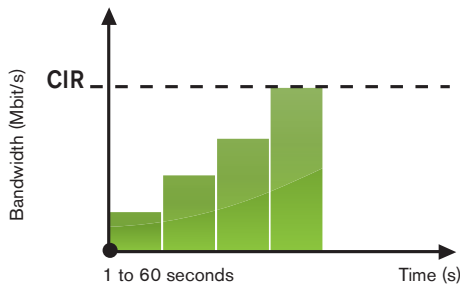
5.1 SERVICE CONFIGURATION TEST

The service configuration test is a per-service test that verifies the bandwidth and performance requirements of a specific service, as defined by the user. The process follows three key phases and monitors all performance indicators during these steps to ensure that they are all met at the same time.

5.1.1 Phase 1: Minimum Data Rate to CIR

In this phase, bandwidth for a specific service is ramped up from a minimum data rate to the committed information rate (CIR). This ensures that the network is able to support this specific service at different data rates while maintaining the performance levels. It also provides a safe and effective way to ramp up utilization without overloading a network in the event that the service is not configured correctly.

As the service is gradually ramping up to the CIR, the system automatically measures KPIs at each step to ensure that the minimum performance objectives are always met. For this phase to pass, all performance objectives must be met at each step all the way up to CIR.



- Service is ramped up from minimum value to CIR rate
- At each step, KPIs are measured and validated against pass/fail criteria

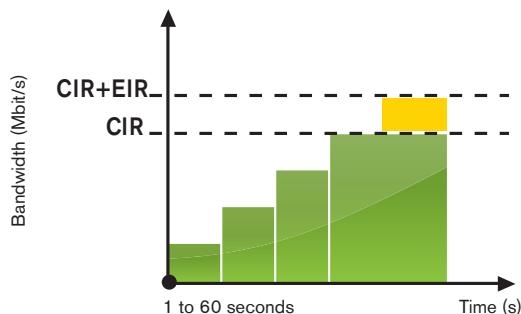
CIR Pass/Fail Criteria:

- ✓ Rx rate = Tx rate
- ✓ KPIs within SLA
- ✗ Rx rate < Tx rate
- ✗ Any KPI fails

5.1.2 Phase 2: CIR to EIR

In this phase, the service is ramped up from the CIR to the excess information rate (EIR). This ensures that the service's EIR is correctly configured, and that the rate can be attained. However, as per accepted principles, performance is not guaranteed in the EIR rates; therefore, no KPI assessment is performed.

At this stage, the system only monitors the received throughput. Since EIR is not guaranteed, bandwidth may not be available for all traffic above the CIR. A pass condition corresponds to the CIR as the minimum received rate and the EIR as a possible maximum. Any measured rate below the CIR is considered as having failed.



- Service is tested at EIR
- KPIs are not guaranteed
- Pass/fail based on Rx throughput

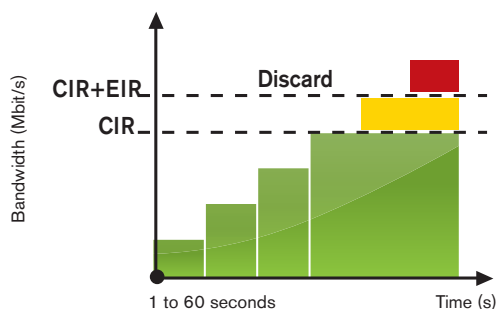
EIR Pass/Fail Criteria:

- ✓ $CIR \leq Rx \text{ rate} \leq CIR+EIR$
- ✗ Rx rate < CIR

5.1.3 Phase 3: Traffic Policing Test

One of the attributes of packet transport is the capability to handle bursty traffic. EIR can occur in conditions of burst, or conditions that surpass the committed bandwidth, and this usually leads to discarded traffic.

In this step, traffic is sent above the EIR, and the received rate is monitored. At the very least, the CIR must be forwarded. The EIR traffic should be forwarded depending on the availability of resources. Any traffic above this maximum should be discarded in order to avoid overloading the network. If the traffic received exceeds the EIR, this means that a device is not properly configured and a fail condition will be signaled.



- Service is tested at traffic policing rate
- KPIs are not guaranteed
- Pass/fail based on Rx throughput

Traffic Policing Pass/Fail Criteria:

- ✓ $CIR \leq Rx \text{ rate} \leq CIR+EIR$
- ✗ Rx rate > CIR+EIR

These three phases are performed per service; therefore, if multiple services exist on the network, each service should be tested sequentially. This ensures that there is no interference from other streams, and that the bandwidth and performance of the service alone are measured specifically.

At the end of the Ethernet service configuration test, the user has a clear assessment of whether the network elements and path have been properly configured to forward the services while meeting minimum KPI performance objectives.

5.1.4 Phase4: Burst Testing

The burst test is a subtest within the service configuration test. In the context of SLA assessments, the objective of burst testing is to verify that the expected burst size can be transmitted through the network equipment with minimal loss. The bandwidth profile of the network equipment contains attributes of committed burst size (CBS) and excess burst size (EBS) that service providers should test at the time of service activation to verify proper attribute configuration.

The most common protocol used today to transport data in IP-based networks is transmission control protocol (TCP). TCP is a bursty protocol by nature. Therefore, it is very useful for service providers to perform burst testing for TCP-based applications such as FTP, HTTP and e-mail services during their service turn-up and troubleshooting phases.

The burst test phase is composed of two parts, the CBS and EBS test. The CBS is the number of allocated bytes available for bursts transmitted at rates above the CIR while meeting the SLA requirements. The EBS is the number of allocated bytes available for bursts transmitted at rates above the CIR+EIR while remaining EIR-conformant. The following graphs show examples of the CBS and EBS burst test sequence.

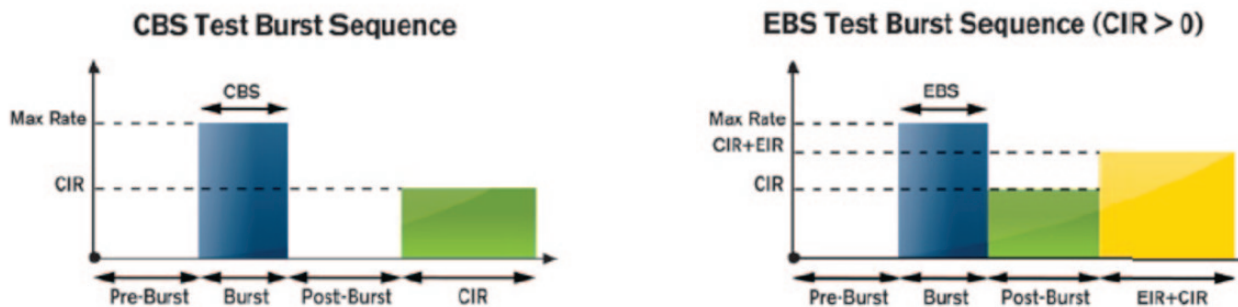


Figure 1. CBS and EBS burst test sequences

Because the CBS and EBS attributes on the network equipment may be configured differently for each service direction, testing CBS and EBS in a round-trip configuration (one end in loopback) has little to no value. It is essential that these parameters be tested independently for each service direction.

Leveraging EXFO's simultaneous bidirectional testing, network operators can stress and emulate real-life network traffic during their service turn-up and troubleshooting phases. This is truly the only way to accurately test and validate proper network configuration and operation—especially when testing with a bursty traffic type such as the previously mentioned TCP.

Along with the CBS and EBS burst size, the burst sequence parameters are fully configurable on EXFO's NetBlazer and PowerBlazer series, as shown in the diagram below.

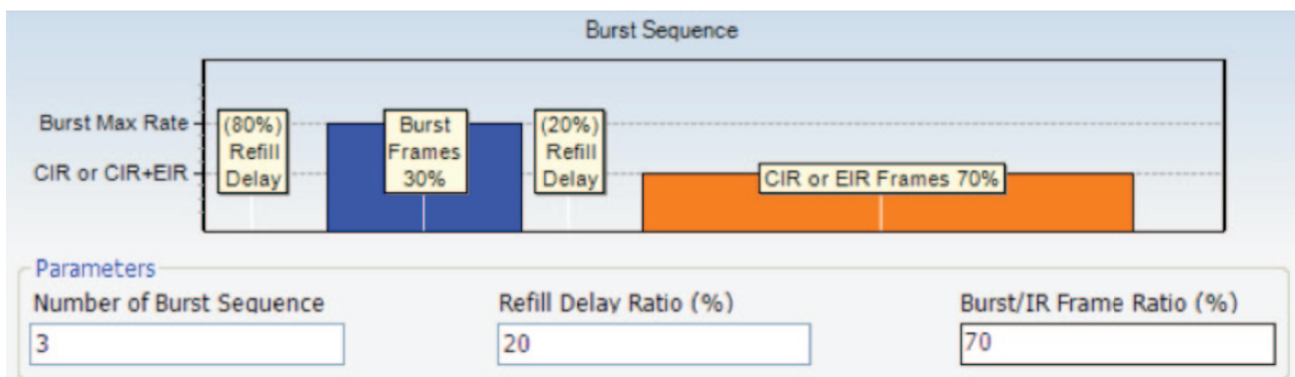
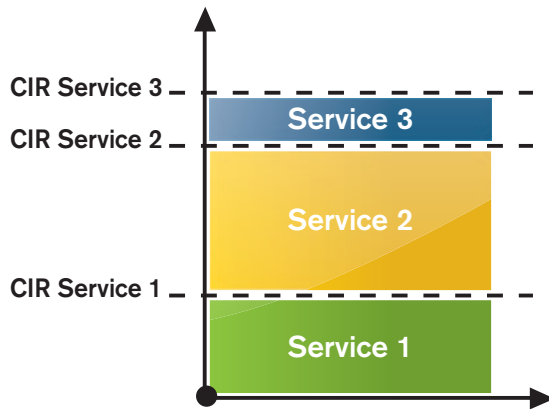


Figure 2. Burst sequence parameter configuration



5.2 SERVICE PERFORMANCE TEST

While the service configuration test concentrates on the proper configuration of each service in the network elements, the service performance test focuses on the enforcement of the QoS parameters under committed conditions, replicating real-life services.

In this test, all configured services are generated at the same time and at the same CIR for a soaking period that can range from a few seconds to a maximum of 30 days. During this period, the performance of each service is individually monitored. If any service fails to meet its performance parameters, a fail condition is signaled.



Service Test Pass/Fail Criteria:

-  • KPIs within SLA per service
-  • Any KPI fails

The combination of these two tests provide all the critical results in a simple and complete test methodology. The service configuration test quickly identifies configuration faults by focusing on each service and how it is handled by the network elements along the paths. The service performance test focuses on the network's capacity to handle and guarantee all services simultaneously. Once both phases have been successfully validated, the circuit is ready to be activated and placed into service.

5.3 ETHERSAM TEST TOPOLOGIES: LOOPBACK AND BIDIRECTIONAL (DUAL TEST SET)

EtherSAM can also perform round-trip measurements with a loopback device. In this case, the measured value reflects the average of both test directions, from the test set to the loopback point and back to the test set. In this scenario, the loopback functionality can be performed by another test instrument in Loopback mode, or by a network interface device (NID) in Loopback mode.

The same test can also be launched in Dual Test Set mode. In this case, two test sets, one designated as local and the other as remote, are used to communicate and independently run tests simultaneously for each direction. This provides much more precise test results, such as independent assessment per direction and the ability to quickly determine which direction of the link is experiencing failure.

It is important to point out that EXFO's EtherSAM test application performs a simultaneous bidirectional test, which means that traffic is active in both directions simultaneously. Testing today's advanced network paths simultaneously in both directions is crucial. This emulates real-life network traffic, and can uncover network equipment configuration issues that could go undetected with non-simultaneous bidirectional testing. Furthermore, performing simultaneous bidirectional testing significantly reduces costs by decreasing test time by 50%.

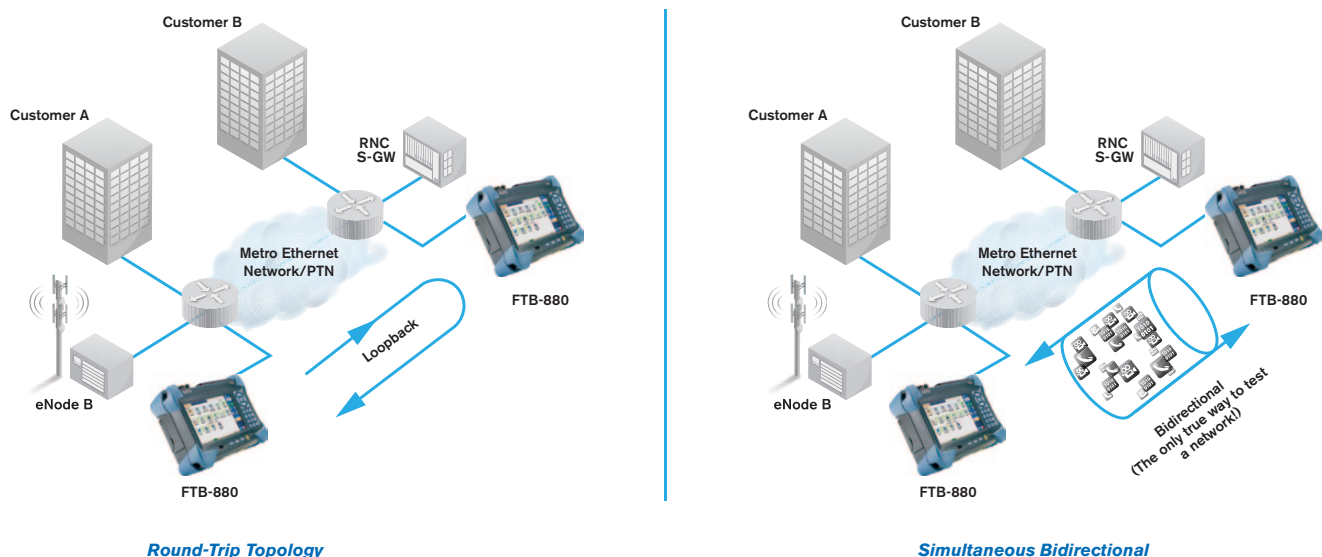


Figure 2. Round-trip topology and simultaneous bidirectional

6 BENEFITS OF ETHERSAM (ITU-T Y.1564)

EtherSAM provides numerous benefits to service providers offering mobile backhaul, commercial and wholesale Ethernet services.

6.1 MEASUREMENT OF ALL KPIs IN A SINGLE TEST

While existing methodologies like RFC 2544 only provide the capability to measure the maximum performances of a link, EtherSAM uses a validation approach in which KPIs are measured and compared to expected values for each service. This approach focuses on proving that KPIs are met while in guaranteed traffic conditions.

KPI	RFC 2544	EtherSAM
Throughput	RFC 2544 makes no distinction between committed and excess traffic, so it always tests the EIR, where performance is not guaranteed by an SLA.	EtherSAM clearly tests the CIR, the EIR and discard traffic conditions, ensuring that the rate-limiting and traffic-shaping functions are correctly configured on network elements.
Frame Loss	Frame loss assessment is based on rate distribution, meaning that frames are generated at specific intervals of transmission rates. This does not fit with the committed and excess profiles, because these key thresholds are not covered by the frame loss distribution.	Frame loss is constantly assessed during the committed phase, ensuring that any loss is quickly identified. This is performed during the throughput test, thus reducing the test time required to validate a service.
Frame Delay	Frame delay is based on the measurement of a single frame during a test time. This approach does not take into account any variation or peak that can occur during a long test time.	EtherSAM measures latency on all generated frames, and provides the peak latency and average latency during the committed range. This ensures that any inconsistency or threshold crossing in the committed range is identified, providing the true latency characteristics of a service.
Frame Delay Variation	RFC 2544 does not measure this KPI. Another test is required.	EtherSAM automatically measures this KPI during the committed phase, ensuring that real-time services are prioritized and forwarded within specifications.

Table 4. RFC 2544 vs. EtherSAM

6.2 SIGNIFICANTLY FASTER

The RFC 2544 methodology uses a sequential approach, where each subtest is executed one after the other until they have all been completed, making it a time-consuming procedure. Additionally, the completion of a subtest relies heavily on the quality of the link. If it experiences many events that cause frame loss, test time dramatically increases. This approach introduces inconsistent delays in the test process, because there is no clear way to determine how long a subtest will last. Furthermore, the RFC 2544 subtest cycles through a distribution of frame sizes, which increases the total test time.

In opposition, EtherSAM uses a defined ramp-up approach in which each step takes an exact amount of time. Because a pass/fail condition is based on KPI assessment during the step, link-quality issues are quickly identified without increasing test time. Each step can be configured from one to sixty seconds, with a default of five seconds. This approach is much more efficient, because it produces valid results in a very short amount of time as compared with the RFC 2544 algorithm.

Based on an equivalent test configuration for a single service, RFC 2544 takes 20 minutes to complete a throughput/latency cycle, whereas EtherSAM takes 25 seconds to complete the entire network configuration test while providing more KPI statistics. This represents a 98% gain in test time.

Here is a comparison of the RFC 2544 throughput methodology and the EtherSAM service performance test for a similar service:

RFC2544

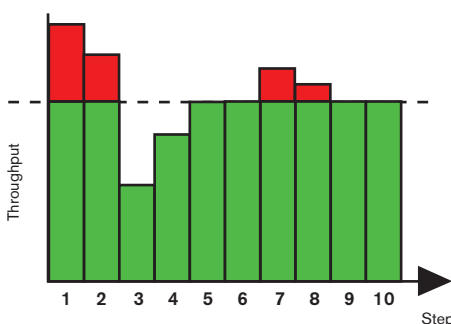
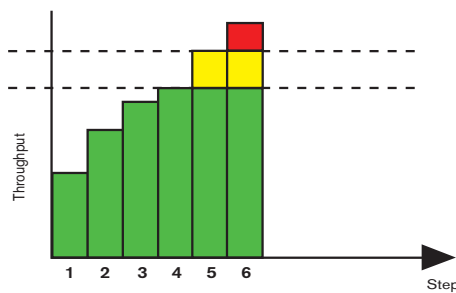


Table 5. Test time for RFC2544

Trial	Tx Rate	Pass/Fail
1	100%	Fail
2	80%	Fail
3	40%	Pass
4	60%	Pass
5	70%	Pass
6	75%	Pass
7	77.5%	Fail
8	76%	Fail
9	75%	Pass
Validation	75%	Pass

EtherSAM



Trial	Tx Rate	Notes
1	50% CIR	Green traffic
2	75% CIR	
3	90% CIR	
4	CIR	
5	EIR	Yellow traffic
6	Traffic policing	Red traffic

Table 6. Test time for EtherSAM

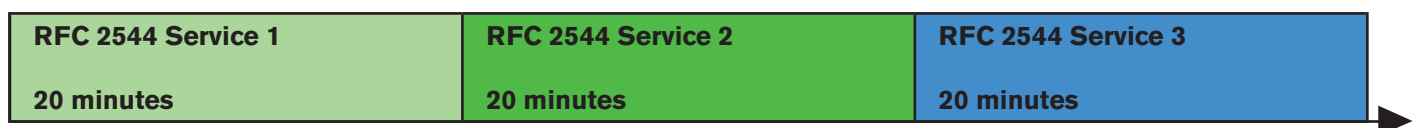
Method	Test Time per Frame Size	Test Time for Seven Standard Frame Sizes
RFC 2544	10 x 5 seconds = 50 seconds	50 x 7 = 5 minutes 30 seconds
EtherSAM	6 x 5 seconds = 30 seconds	30 x 7 = 3 minutes 30 seconds

6.3 MULTISERVICE CAPABILITIES

As described earlier, the majority of Ethernet services deployed today include multiple classes of service within the same connection. For example, 3G mobile backhaul services typically comprise four classes of service, while 4G LTE backhaul has seven. A major drawback of the RFC 2544 is the fact that it can only test a single service at a time. This implies that the RFC 2544 test must be relaunched for all services during a configuration test. Combining the need to repeat the test per distribution with the time it takes to complete a single cycle, RFC 2544 inefficiently increases test time.

In comparison, EtherSAM individually cycles through the different services, providing quick assessment of the service configuration. This dramatically reduces the total test time of multiple services. In a typical three-service performance test scenario, the total RFC 2544 test takes around one hour, while the same assessment with the same frame size using EtherSAM takes nine minutes.

Three Service RFC2544 Test



Three service EtherSAM Test



Method	Test Time per Service	Test Time for three services
RFC 2544	20 minutes per service	60 minutes for three services
EtherSAM Service Configuration	3 minutes per service	9 minutes for three services

Table 7. Three service EtherSAM tests

6.4 MORE REPRESENTATIVE OF REAL-LIFE CONDITIONS

The worst-case scenario for a network is handling multiple traffic types during a congestion period. During these high-stress conditions, network equipment must perform prioritization in accordance with the configured SLA.

Since RFC 2544 only tests one stream at a time and each service individually, it cannot simulate worst-case scenarios.

The EtherSAM service subtest can generate all configured services at the same time, providing the ability to stress network elements and data paths in worst-case conditions. The service performance test provides powerful test results, since all KPIs are measured simultaneously for all services with clear pass/fail indication and identification of failed KPIs. This ensures that any failure or inconsistency is quickly pinpointed and reported, again contributing to an efficient and more meaningful test cycle.

7 ETHERSAM TEST SCENARIOS

EtherSAM can be used for a number of test applications:

Test Application	Purpose	Benefits of the EtherSAM Methodology
Lab/Production	<ul style="list-style-type: none"> › Testing new equipment before deployment › Testing major configuration changes 	<ul style="list-style-type: none"> › The network configuration test quickly determines if the new equipment or the configuration change is performing as expected for each service. › The service performance test provides full load analysis by testing multiple services simultaneously, and ensuring that all services meet their guaranteed performance objectives.
Service Turn-Up	<ul style="list-style-type: none"> › Validating the SLA before turning up a service › Issuing a birth certificate before delivering a circuit to a customer › Gathering performance metrics for traceability 	<ul style="list-style-type: none"> › EtherSAM ensures that services are correctly configured before turn-up, when the available testing window is at its minimum. › Subtests can be used for SLA verification and proof of performance for customers. › A detailed test report ensures the traceability of performance results in the event of trouble tickets; it also serves as a birth certificate for new circuits.
Troubleshooting	<ul style="list-style-type: none"> › Resolving service tickets or ensuring that customer complaints are received › Fixing degradations identified by the monitoring system 	<ul style="list-style-type: none"> › EtherSAM can be used as a troubleshooting tool to quickly identify which service/KPI is failing on a link. › EtherSAM is available on fixed and portable test instruments, providing field technicians with more flexibility. › Technicians can test from any location. › Once a fix is completed, EtherSAM can be used to validate the new configuration and ensure that existing services have been maintained.

Table 8. EtherSAM test applications

7.1 LOOPBACK TEST SCENARIO

The simplest EtherSAM test scenario is the loopback test, where EtherSAM testing is performed from test devices to loopback points placed at key locations. The loopback ecosystem includes three elements.

7.1.1 Centralized Test Points

EXFO's BV-3100 in-network, multifunction service assurance verifier, are typically located at fixed test points, such as in the central office and data centers. They provide centralized access to integrated turn-up, monitoring and troubleshooting capabilities.

7.1.2 Portable Test Equipments

Portable test devices bring powerful testing capabilities directly to the field and customer sites. These devices include:

- › FTB-8510B and FTB-860 Gigabit Ethernet Test Module
- › FTB-8510G, FTB-860G and FTB-880 and 10 Gigabit Ethernet Test Modules
- › FTB-8130NGE and FTB-88100NGE Power Blazer Next-Generation Multiservice Test Module

7.1.3 Loopback Devices

Loopback devices offer responder capabilities to complete round-trip testing. These devices, which provide wire-speed capabilities and perform address swapping to maintain integrity in the network, include:

- › Dedicated loopback points such as EXFO's ETS-1000L Ethernet Loopback Device
- › Portable devices with loopback capabilities up to 10 Gbit/s such as EXFO's FTB-860GL Ethernet Test Set
- › Dedicated responders such as EXFO's BV10 Performance Endpoint Unit

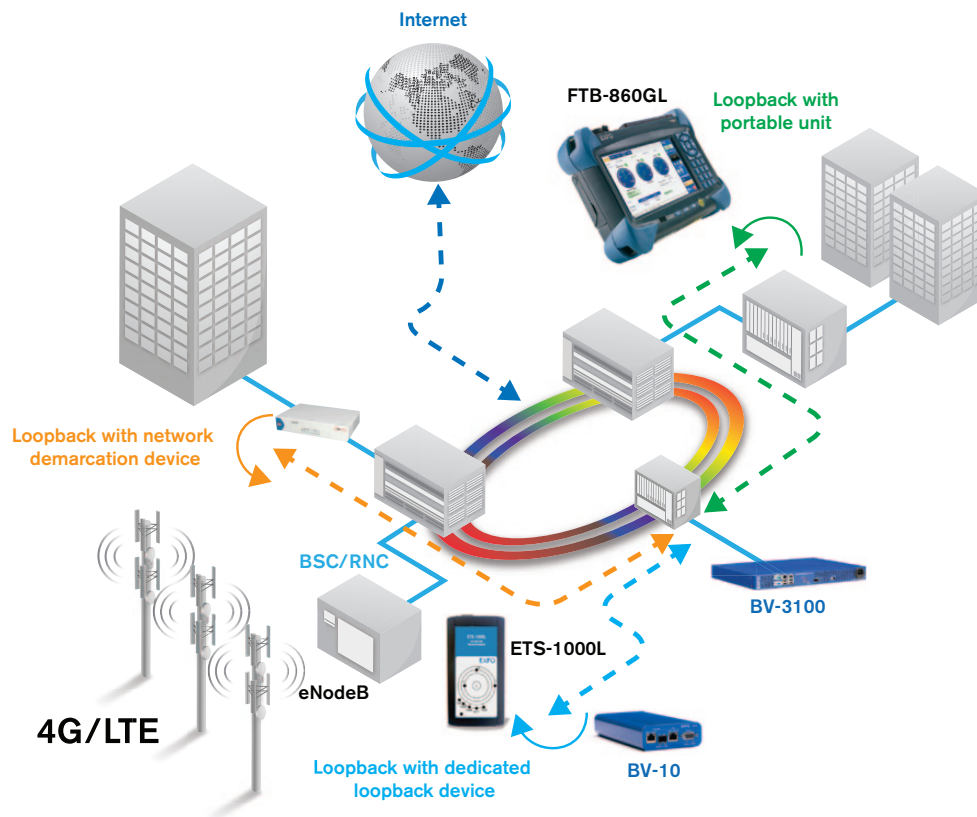


Figure 3. Testing with a loopback device

7.2 BIDIRECTIONAL (DUAL TEST SET) TEST SCENARIO

The dual test-set approach is a more accurate test scenario. In this case, two units perform an asymmetrical SLA measurement, providing EtherSAM results per direction. This scenario's main strength is that it quickly pinpoints which direction has not been configured properly or is at fault, while providing performance metrics per direction.

Results from both directions are sent and displayed on the local unit. This ensures that the entire test routine can be completed by a single technician in control on a single unit, providing shorter test times and reduced manpower. This flexibility also guarantees that different units can be set as remote units. The most interesting scenario is a centralized unit that is always configured as a remote unit with fixed addresses. The carrier can simply dispatch one field technician to the test site. The tester will then quickly and efficiently execute service turn-up and burn-in without requiring an extra worker at the central office.

The dual test-set approach also provides the capability needed to segment the network and quickly pinpoint the direction in which issues occur. This is especially important where bandwidth is different between the upstream and downstream direction. In such cases, using a loopback tool will always yield the same results. The measurement will be affected by the lowest throughput, and the test results will not show that one direction has a higher performance than the other. The same scenario will occur if a network misconfiguration is present in only one direction of the service. Depending on the error, the problem will not be identified with round-trip measurements. This often results in customer complaints and additional truck rolls.

With the dual test-set approach, both directions are independently analyzed simultaneously and pass/fail results are provided per direction, yielding the highest level of confidence in service testing at the lowest cost.

This test scenario requires two test units, which can include:

- › Centralized test points: These instruments, such as EXFO's RTU-310/310G IP Services Test Heads, are typically located at fixed test points such as in central offices and data centers. These instruments provide centralized access to test capabilities.

- › Portable test devices bring powerful testing capabilities directly to the field and customer sites. These devices include:

FTB-8510B and FTB-860 Gigabit Ethernet Test Module

FTB-8510G, FTB-860G and FTB-880 and 10 Gigabit Ethernet Test Modules

FTB-8130NGE and FTB-88100NGE Power Blazer Next-Generation Multiservice Test Module

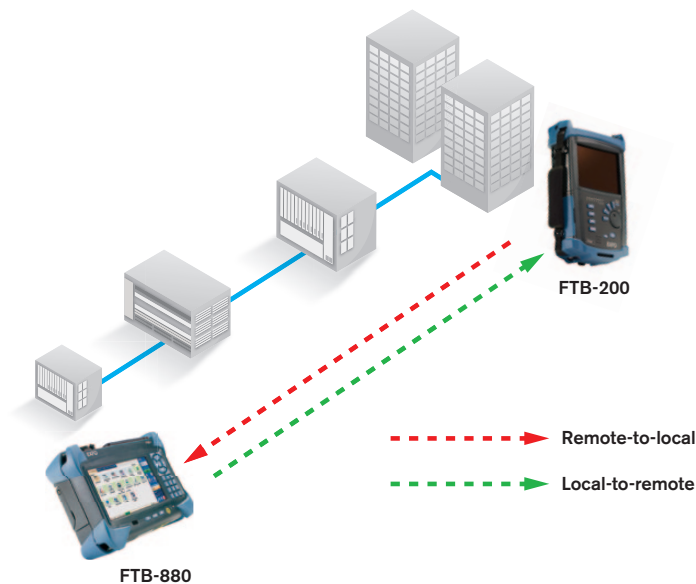


Figure 4. Testing in both directions simultaneously

CONCLUSION

Carriers and service providers face the constant challenge of ensuring the proper delivery of services to customers. Ethernet services need to be delivered to customers in compressed time frames, while proving to be more reliable than ever. The original method of assessing performance via RFC 2544 no longer provides them with the tools needed to turn-up and troubleshoot the services required by customers with a high level of confidence.

EtherSAM bridges the gap between service validation and performance assessment, by providing an intuitive and easy approach to confidently control and manage networks while reducing OPEX and growing revenues. EtherSAM is the only standard test methodology that allows for complete validation of SLAs in a single, significantly faster test, while offering the highest level of accuracy.

EXFO Headquarters > Tel.: +1 418 683-0211 | Toll-free: +1 800 663-3936 (USA and Canada) | Fax: +1 418 683-2170 | info@EXFO.com | www.EXFO.com

EXFO serves over 2000 customers in more than 100 countries. To find your local office contact details, please go to www.EXFO.com/contact.